



(19)

Europäisches Patentamt

European Patent Office

Office européen des brevets



(11)

EP 0 961 518 A2

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
01.12.1999 Bulletin 1999/48(51) Int. Cl.⁶: H04Q 11/04, H04L 12/56

(21) Application number: 99109317.0

(22) Date of filing: 28.05.1999

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE
Designated Extension States:
AL LT LV MK RO SI(72) Inventors:
• McAllister, Shawn
Nepean, Ontario K2G 5V9 (CA)
• Tooker, Mark
Kanata, Ontario K2L 2C7 (CA)
• Veeneman, Ron
Carleton Place, Ontario K7C 4E4 (CA)

(30) Priority: 28.05.1998 CA 2239032

(71) Applicant:
Newbridge Networks Corporation
Kanata, Ontario K2K 2E6 (CA)(74) Representative:
Modiano, Guido, Dr.-Ing. et al
Modiano, Josif, Pisanty & Staub,
Baaderstrasse 3
80469 München (DE)

(54) Operator directed routing of connections in a digital communications network

(57) A virtual circuit whose path across a network is specified by a human operator (hereinafter operator directed route or ODR VC) is established by manually provisioning a preferred path for the connection, including a source node, destination node, and intermediate nodes or subnetworks therebetween. The source node creates a call setup message which is signalled along the preferred path, whereby the intermediate nodes along the preferred path establish the bearer channel cross-connects. The operator also specifies a re-routing scheme for the ODR VC in the event the preferred path

is blocked or a link subsequently fails. The re-routing scheme includes one of (a) the preferred path only; (b) at least one manually provisioned alternate path; and (c) any available path. The ODR VC provides the benefit of a permanent virtual circuit (PVC) in terms of the ability to consciously route a connection with the benefit of a soft permanent virtual circuit (SPVC) or switched virtual circuit (SVC) in terms of the capability to efficiently re-route connections by the network as opposed to a central management authority.




Figure 1

EP 0 961 518 A2

Description

Field of the Invention

[0001] The invention generally relates to the art of establishing connections in a network; and more specifically to the establishment of a virtual circuit in an asynchronous transfer mode (ATM) network using an operator-directed routing path.

Background of the Invention

[0002] A permanent virtual circuit (PVC) provides a bearer channel path across a network which comprises a series of bearer channel links that are interconnected by "permanent" bearer channel cross-connections established on network nodes under the direction of a central network management authority. This authority can be a human operator which decides the route and manually configures each cross-connection individually through a network management terminal interface (NMTI). Alternatively, the authority can be a network management system (NMS), which automatically selects the route through the network according to some algorithm or objective when requested by one or more human operators. The NMS is connected to each network node typically through an independent control channel and thereby automatically establishes the cross-connections.

[0003] For each PVC thus established, the central network management authority can choose a route through the network which meets criteria based on network-wide policies rather than single-element policies. One example of a network-wide routing policy is a policy of using the most efficient route through the network, for example, by minimizing the number of network nodes traversed; or by minimizing the cumulative costs of links which are traversed; or, by balancing the number of traversed nodes and the cost of the traversed links. Another example is balancing network usage across different network nodes, or the links between network nodes, such that no one network node or link is carrying a large proportion of all PVCs traversing the network.

[0004] This approach to establishing connections provides various benefits. For instance, since there is only one central network management authority, i.e. an expert human network operator or a powerful computer running sophisticated network management software, the cost of provisioning the authority is inexpensive relative to the overall cost of the network. In addition, whenever network management policy is changed with respect to the routing of circuits or connections, it is easy to implement the policy changes because they only need to be made in one place, in the central management authority.

[0005] This approach also has various shortcomings. One shortcoming is the relatively high cost of maintaining an exact and up-to-date picture of network condi-

tions in the central network management authority in order to enable routing and re-routing decisions to be made with accuracy. Another shortcoming is the slow speed at which the central network management authority can re-route PVCs in the event of a network failure. This is due to the time required for the central authority to (i) become aware of the network failure, (ii) find new routes through the network for all affected PVCs which satisfy all of the various network element and network-wide criteria, and (iii) re-establish all affected PVCs along the chosen routes.

[0006] A soft permanent virtual circuit (SPVC) provides a bearer channel path across a network which comprises a series of bearer channel links that are interconnected through "switched" (i.e. on-demand) bearer channel cross-connections made across a series of network nodes. More specifically, the ingress and egress network nodes are provisioned by an operator (either through the NMTI or NMI) but the cross-connections are commissioned via signalling, like a switched virtual connection (SVC), as the SPVC is signalled and routed across the network from the ingress network node to the egress network node. An SVC is a path that is signalled from user side UNI to user side UNI whereby the route is chosen by the network nodes as the path is signalled from a source and station towards a destination end station. The individual cross-connections for the SVC path are configured and connected by call control software running on each node along the path as the path steers itself through the network using routing tables resident on each node (i.e., hop-by-hop routing) or according to a predetermined route specified in the connection request (i.e., source-routing). Thus, SPVCs are a kind of hybrid between PVCs and SVCs since SPVCs, like PVCs, are initiated by the central network management authority and require no UNI signalling between the user and the network, but, like SVCs, the cross-connections are routed through the network and maintained by the network nodes themselves.

[0007] One of the benefits of this approach to establishing connections is that SPVCs can be re-routed more efficiently because the network nodes which are closest to a network failure can quickly detect the failure and initiate the re-routing procedures. Hence, the virtual circuits can be re-established more quickly and at less cost than PVCs can be re-established by a central network management authority. Indeed, it is estimated that SPVCs improve the fault restoration re-route performance for connections by an order of magnitude because the processing is distributed within the network rather than being centrally managed by the NMS.

[0008] There are also various disadvantages with this approach. First, it is not possible for an expert human operator to intervene and influence the network nodes to use routes for an SPVC through the network that differ from the routes that would be automatically chosen by the network nodes, or to modify routes chosen automatically by the network nodes, for example, in order to

impose routing criteria for which the network nodes have not been designed or configured. Second, due to the large number of network nodes, it is difficult and costly to employ powerful computing devices for each network node; therefore, the sophistication of the routing algorithms implemented by the network elements cannot, at reasonable expense, approach the sophistication of the routing algorithms that can be implemented by a central network management authority which comprises of a smaller number of computing devices. Third, again due to the large number of network nodes, it is difficult and costly to upgrade or reconfigure each network node whenever a network-wide policy has changed with respect to the routing and re-routing of virtual circuits, for example, when a new virtual network has been created out of available resources on many different nodes across the network, or when network management policies change with respect to the weighting of different criteria such as the number of nodes traversed versus cumulative cost of links traversed.

[0009] To further elaborate upon the disadvantages provided by both PVCs and SPVCs, consider for example, the reference network shown in Fig. 1. A customer wishes to connect end station or customer premise equipment (CPE) 20 from Toronto to Montreal, and purchases two (2) connections or virtual circuits 24a and 24b therebetween to ensure redundancy. If the connections 24a and 24b between Montreal and Toronto are established using SPVCs, then, since the path, i.e. intermediate nodes 28 and links 26, of the connections are not preconfigured, a situation such as illustrated in Fig. 1 could occur, wherein the network nodes select the shortest path between the CPEs. Thus, the paths of both connections 24a and 24b are identical, traversing node A, link 26ag, node G, link 26dg, and node D. This result would thus destroy the sought after redundancy. In order to ensure that the path of each connection does not follow a common link or share the same physical interface of the other, it is possible to provision the connections as PVCs in order to manually configure the cross-connections and predetermine the followed links. However, this strategy brings with it the above described disadvantages of PVCs, in particular, the relatively slow re-route performance of the centralized NMS in the event of a service disruption such as a failed link. For customers who have demanding maximum permissible service disruption requirements, e.g., one second per year, the re-route performance of PVCs by the NMS is unacceptable. For example, it may take the NMS two hundred (200) seconds to re-route a severed OC-3 cable carrying ten thousand (10,000) connections at fifty (50) re-routes per second.

[0010] Similar problems arise in router based network environments such as those networks which employ tag switching or multi-protocol label switching (MPLS).

Summary of the Invention

[0011] The invention seeks to avoid various limitations of the prior art by providing a hybrid type of connection which features the fast, distributed processing re-route capabilities of SVCs or SPVCs, yet enables a human operator to direct the routing of the path across the network like a PVC. Broadly speaking, one aspect of the invention relates to a communications network comprising a user interface means for enabling a human operator to manually provision a predetermined preferred path across the network, which path comprises a source network element, a destination network element, as well as intermediate network elements and, optionally, the intermediate ports or links between these network elements. The user interface means is connected to the source network element which encapsulates the manually provisioned preferred path in a source-routed connection request message that is signalled to the network elements along the preferred path. The network elements include call processing means to ensure that the manually routed connection traverses these intermediate nodes and the specified intermediate links as the connection request is signalled from the source network element to the destination network element.

[0012] Another aspect of the invention relates to a method of establishing a connection in a communications network comprising interconnected network elements. The method comprises:

- (a) providing a predetermined preferred path definition to a source network element, wherein the path definition comprises a destination network element and intermediate network elements between the source and destination network elements;
- (b) creating a source-routed connection request message which specifies the preferred path; and
- (c) signalling the connection request message from the source network element to the destination network element along the preferred path in order to establish the connection.

[0013] The network elements may be end stations, network nodes, or entire subnetworks. In the preferred embodiment the network elements are network nodes, and the manually routed connection resembles an SPVC. This type of connection is henceforth referred to as an "operator directed route" or "ODR" SPVC.

[0014] In the preferred embodiment, the network is a P-NNI network, as described in greater detail below, and the source routed connection request is an SPVC call setup message which carries a designated transit list (DTL) specifying the preferred path for routing the ODR SPVC connection.

[0015] In the preferred embodiment, the operator provisions not only the preferred or primary path, but also

one or more alternate paths for the ODR SPVC. The primary path comprises the nodes and specified links an operator prefers the ODR SPVC to traverse as it is routed across the network. The alternate paths comprise alternate nodes and links for the ODR SPVC, in the event the primary path is inoperative for any reason and must be re-routed. An operator preferably also specifies re-route restrictions, that is, whether the ODR SPVC is restricted to traversing the network along the specified primary and/or alternate paths whenever the ODR SPVC is re-routed, or whether any available path may be traversed.

[0016] ODR SPVCs, having paths provisioned by a human operator, offer the same services as the prior art SPVCs, and provide various other additional benefits. For instance, ODR SPVCs allow the operator to deliberately control the distribution of connections across a network. Additionally, the operator can include or exclude certain nodes(s) from the path of an ODR SPVC for security reasons. For example, a connection may need to traverse certain nodes to ensure the connection is protected from unauthorized access. This offers some degree of information security for the customer. In addition, the operator can setup ODR SPVCs to ensure a network failure will not disrupt all of a customer's connections by requiring a certain portion of the customers' connections to traverse a certain path while the remainder traverse completely unrelated path(s). If any path(s) experiences failure, the failure would not disrupt all of the customer's connections. Furthermore, ODR SPVCs can be provisioned across a network to ensure that the network is efficient. For example, the operator can direct certain ODR SPVCs to traverse the network over the least amount of hops, thereby guaranteeing that a connection will take the most direct path to its destination. The operator may even direct ODR SPVCs to avoid certain nodes to free up resources and avoid congestion in those nodes.

Brief Description of Drawings

[0017] The foregoing and other aspects of the invention will become more apparent from the following description of the preferred embodiment thereof and the accompanying drawings which illustrate, by way of example, the preferred embodiment. In the drawings:

- Fig. 1 is a diagram of a reference network;
- Fig. 2A is a diagram of an P-NNI reference network having the same topology as the network shown in Fig. 1;
- Fig. 2B is a diagram of the P-NNI reference network illustrating desired ODR SPVC paths;
- Fig. 3 is a diagram illustrating the logical partitioning of the bandwidth provided by a physical interface, in accordance with the preferred embodiment;
- Fig. 4 is a diagram of a P-NNI control plane; and
- Fig. 5 is a database relationship diagram illustrating

how an ODR SPVC is represented in a database on a source node in accordance with the preferred embodiment.

Detailed Description of Preferred Embodiments

[0018] Fig. 2A illustrates a reference connection-orientated ATM network 30 having a similar topology to the reference network shown in Fig. 1. The network 30 comprises a plurality of interconnected network elements, including end stations 20 and network nodes 32. For ease of reference, individual nodes are identified by an alphabetical suffix, e.g., A, B or C and referenced elements of a given node are also generally labelled with the same suffix used to identify the node. The nodes 32 include various ports (not shown) which are interconnected through standard physical interfaces 34, such as well known OC-3, OC-12 and DS3 fibre optic or electrical interfaces.

[0019] In the preferred embodiment, the nodes 32 interface with one another using the Private Network-to-Network Interface (P-NNI) protocol which is described in the reference "Private Network-to-Network Interface Specification Version 1.0 (P-NNI 1.0)", doc. no. af-p-nni-0055.00, March 1998, published by the ATM Forum, and which is incorporated herein by reference in its entirety. The nodes 32 are interconnected by data links 36, each of which represents a pre-allocated portion of the bandwidth provided by the corresponding physical interface 34. (Note that multiple links can be associated with each physical interface spanning adjacent nodes.) In the preferred embodiment, each data link 36 is associated with a P-NNI signalling link 38 and a P-NNI routing link 40, alternatively termed routing control channel (RCC), which span adjacent nodes. Fig. 3 shows more accurately the relationship between a physical port or interface, data links 36 (which are alternatively referred to as "trunk groups"), P-NNI signalling links 38, P-NNI routing links 40, SVCs and SPVCs. It will be seen from the foregoing that the data link 36 represents a P-NNI connectivity between two nodes 32.

[0020] More specifically, each node 32 comprises a P-NNI signalling module 42 for carrying out a P-NNI signalling protocol, which is based on ATM Forum UNI signalling with extensions to support P-NNI functions. Signalling module 42 communicates over the P-NNI signalling link 38, which may be a designated PVC or SVC associated with the data link 36. Collectively, the signalling modules 42 of the network nodes and associated signalling links 38 therebetween provide a signalling network for forwarding or carrying P-NNI signalling protocol messages, including connection-orientated messages such as the Setup, Connect and Release messages defined in the P-NNI reference, to and between network nodes 32. Similarly, each network node 32 comprises a P-NNI routing module 44 for carrying out a P-NNI routing protocol wherein nodes exchange topology information with one another over

the P-NNI routing links 40 in order to dynamically compute paths through the network. The P-NNI routing links 40, which also may be PVCs or SVCs associated with the corresponding data link 36, carry P-NNI routing protocol messages such as Hello, PTSP, Database Summary, PTSE request and PTSPs acknowledgement messages to and from neighbour nodes. Further information regarding the P-NNI routing and signalling protocols may be found in the above-noted P-NNI reference.

[0021] The network 30 is also connected to a centralized network management system (NMS) 46 such as described above. In the preferred embodiment, the NMS 46 is connected to the network nodes 32 through a virtual control channel 47 (not entirely shown), but in alternative embodiments the NMS 46 can be connected to the nodes 32 through an overlay control network, e.g., the public telephone network.

[0022] The objective of the preferred embodiment is to establish or set up a connection between end stations or customer premise equipment (CPE) 20A and 20B, which may not support signalling, through the network 30. As previously described, the NMS 46 enables a human operator to establish a prior art PVC or SPVC. In the preferred embodiment, however, the NMS 46 merely provides a user interface means to enable a human operator to manually provision the ODR SPVC of the invention. The ODR SPVC according to the preferred embodiment comprises at least two attributes: (a) a manually provisioned preferred or primary path for the connection; and (b) a manually provisioned re-routing restriction or scheme for re-routing the connection in the event the preferred or primary path is blocked or a data link 36 thereof fails therein. In the preferred embodiment the re-routing scheme may, as explained in greater detail below, includes one or more alternate preferred paths which must also be manually provisioned by the operator.

[0023] Each path which is manually provisioned by the operator comprises a list of all of the nodes and, optionally, the operator-specified links, the ODR SPVC must traverse to reach its destination endpoint, i.e., the egress node.

[0024] The network 30 is associated with an addressing scheme, and thus each node is identified by a unique node identifier, and each data link is represented by a link identifier, which may be local to the node. The format of the link identifier varies depending upon the specific scheme chosen by the manufacturer of the node, but typically represents some sort of numerical identifier with respect to a given node.

[0025] For example, referring to Fig. 2B, in order to ensure redundancy it may be desired to provide one ODR SPVC connection 48 between CPE 20A and CPE 20B through path [node 32A; node 32G; and node 32D], and another ODR SPVC connection 49 between CPE 20A and CPE 20B through path [node 32A; node 32B, link 36BC; node 32C, node 32D]. Thus, the operator uses the NMS 46 to select or confirm the appropriate

node and link identifiers for these paths. The operator may also specify or confirm alternate preferred paths for these connections in the event the preferred paths or routes are inoperative. For example, the operator could specify an alternate route for connection 48 as [node 32A, node 32E, node 32F, and node 32D]. Similarly, the alternate route for connection 49 could be [node 32A, node 32B, link 36BC, node 32C; node 32D].

[0026] The user interface means provided by the NMS facilitates the input of the ODR SPVC, including the preferred or primary path, the optional alternate path, and the re-routing scheme thereof. The user interface means for specifying paths can comprise any of the following methods:

(1) **Manual Entry.** Here node identifiers and optional link identifiers are manually typed into a terminal separated by predefined demarcation symbols.

(2) **Manual with Point-and-Click Assistance.** The NMS provides a graphical representation of the network as known in the art *per se*, and the operator performs "paste" operations to create a list of nodes and optional links for the ODR SPVC path. The actual node and link identifiers are typically hidden from view and replaced by mnemonics which are more readily understood by the operator; nevertheless, the resulting list is specified in the terms of the node identifiers and optional link identifiers.

(3) **Automatic Route Generation with Optional Manual Editing.** The NMS, in a manner similar to the production of a prior art PVC, can automatically generate a path based on predetermined algorithms, such as the shortest path or least cost. This path is displayed, either textually or graphically, to the operator, who may then confirm or edit the path chosen by the software running on the NMS.

[0027] Similarly, the operator also selects a re-routing restriction or scheme through the user interface means, thereby indicating how strictly an ODR SPVC is restricted to routes along the nodes and links provisioned by the operator when circumstances dictate the ODR SPVC must be re-routed. The preferred re-routing schemes, which are discussed in greater detail below, include:

- (a) primary path;
- (b) primary path-alternate path;
- (c) primary path-any path; and
- (d) primary path-alternate path-any path.

[0028] The user interface means may alternatively be provided through the NMTI of a source node, although such interfaces tend to be only textually based. Irrespective of whether the user-interface means is pro-

vided by the NMS or NMTI, the particular user interface methods specified above are all instances of manual provisioning.

[0029] Once the operator has manually provisioned the preferred or primary path, the optional alternate path(s), and the re-routing scheme for the ODR SPVC (i.e. the configuration data), the NMS 46 sends the configuration data and a message over the control channel 47 instructing the ingress node 32A to set up or establish the ODR SPVC. The ingress node 32A stores the ODR SPVC configuration data as described in greater detail below. The signalling module of the ingress node 32A creates a call establishment or connection request message, such as an SPVC Call Setup message specified in Annex C of the P-NNI reference. However, unlike the situation corresponding to the establishment of a prior art SPVC, the ingress node 32A does not automatically compute the designated transit list (DTL) (specifying the source-routed path of an SVC or SPVC) using its P-NNI routing tables. Instead, the ingress node creates a DTL using the ODR SPVC primary path provisioned by the operator. The DTL is included as an information element (IE) in the SPVC Call Setup message.

[0030] The SPVC Call Setup message is signalled by the ingress node over the signalling link to the next or following node in the DTL, which in turn forwards the SPVC Call Setup message to a succeeding node listed in the DTL. This process continues until the egress or destination node receives the SPVC Call Setup message. Each node which receives the SPVC Call Setup message proceeds in the conventional manner to establish or commission a bearer channel cross-connection 51 between an incoming link 34 and an outgoing link 34.

[0031] If a link between two nodes has been provisioned by the operator and thus is specified in the DTL, then, subject to local connection admission control (CAC) verification, the nodes establish the bearer channel cross-connections 51 so as to use the specified link or trunk group, and call processing means 50 operating on the nodes (which comprises the signalling module 42 and routing module 44) needs to only select a free VPI and optionally a VCI on the specified data link or trunk group. If, however, the DTL does not specify any data link or trunk group between nodes, then a local routing function in the call processing means 50 selects an available data link or trunk group to carry the ODR SPVC bearer channel thereon. In either event, it will be seen that since the preferred embodiment uses source-routing, as opposed to hop-by-hop routing, to establish a connection, the path assumed by the connection matches the primary or preferred path specified by the operator, provided that path is not blocked, as discussed in greater detail below.

[0032] Once the destination or egress node 32D receives the SPVC Call Setup message, it returns an acknowledgement message back over the signalling network to inform the ingress or source node 32A that

the ODR SPVC connection has been successfully established. The source or ingress node 32A then sends a message back over the control channel 49 to the NMS 46 which informs the human operator that the ODR SPVC has been successfully established. Thereafter, the operator may configure the CPE 20A and 20B and ingress and egress network nodes as known in the art *per se* to transmit data over the recently established ODR SPVC. It will be seen that since the ingress or source node 32A stores the configuration data for and initiates the ODR SPVC, the nature of the ODR SPVC is transparent to the remainder of the network 30 and it appears thereto that a conventional SPVC has been requested and established.

[0033] The re-route restrictions associated with the ODR SPVC are enforced whenever the network 30 detects that the preferred or primary path of the ODR SPVC is inoperative. This may occur when (a) the ODR SPVC is being initially established and the call setup is blocked along the preferred path, or (b) once the ODR SPVC is up and running and a link failure occurs.

[0034] In the first case, it is possible, for example, that the CAC processing of an intermediate node in the primary path of the ODR SPVC reports insufficient node resources to progress the ODR SPVC through the intermediate node or a specified link thereon in which case the primary path is blocked. In such circumstances, the P-NNI signalling protocol provides a crankback procedure wherein a connection request which is blocked along a selected path is rolled back to a prior node, such as a DTL-creating node (which is node 32A in the reference network). In order to compute another path to an SPVC destination endpoint. The P-NNI crankback procedure is initiated by the network node which detects the block. This node sends a signal indicative of a blocked path, such as a connection clearing message having a crankback IE, back to the source node over the signalling network. The protocol allows a DTL-originating node, such as the source node 32A, which receives this message to re-route the connection using a different path.

[0035] In the second case, the failure of a link in the preferred path of the ODR SPVC may be detected by the normal operation of the P-NNI links 38 and 40. For example, the P-NNI signalling protocol uses the services of the UNI signalling ATM adaption layer 2 (SAAL) (layer 2 in the P-NNI Control Plane shown in Fig. 4) for ensuring a reliable signalling link 38 between nodes. The SAAL protocol constantly sends some protocol messages or data units (PDUs) to a node on the other side of the link which must be acknowledged by the other node. These messages may be sequenced data PDUs, i.e. PDUs which carry layer 3 P-NNI signalling messages for all virtual connections associated with a data link 36, or separately sequenced "poll" and "stat" PDUs. In combination, these messages implement a keep-alive or heartbeat polling process associated with each successive pair of nodes along the path of a call or

virtual connection in order to inform each node in the path that the links to neighbour nodes are alive and functioning. Furthermore, the P-NNI routing protocol employs mechanisms such as flooding, sequence number exchange, "lock step" acknowledgements and check sums to ensure reliable and timely delivery of P-NNI topology state packets (PTSPs). When a link failure is detected, the P-NNI protocols dictate that the functioning part of the network transmit a signal indicative of the failure, such as one of the P-NNI connection clearing messages, to the source or ingress node, specifying the cause for the release. Upon receipt of this signal, the source node may attempt to re-route the connection along a different path.

[0036] The following illustrates recovery operation, whether invoked during call setup caused by a blocked path or during call tear-down caused by a link failure. In the preferred embodiment, when the source node 32A receives a signal indicative of the blocked path or link failure, such as provided by one of the P-NNI connection clearing messages, the source node first determines whether the blocked path is associated with ODR SPVCs managed by the node. If so, the source node attempts to re-initiate the establishment of each such ODR SPVC in accordance with its associated re-routing restriction, as follows:

• Primary Path

[0037] Under the primary path re-routing scheme, the ODR SPVC is restricted to being re-routed only along the nodes and links in its primary path. Failure to route the ODR SPVC along this path will result in the crankback of the connection to its source node. At this point the source node either immediately attempts to re-route, i.e., re-initiate the connect request of, the ODR SPVC along the nodes and links in its primary path or waits an interval of time before attempting to re-establish the ODR SPVC.

[0038] If the failure is due to a link not specified by the operator, the source node immediately re-routes the ODR SPVC along the nodes and links in the primary path. This is done to give the network nodes another opportunity to immediately set up the connection along the nodes and links in the primary path whenever parallel links exist in the area of the link failure. If parallel links exist in the area of the link failure, the appropriate node's call processing means selects a link other than the failed link. The node attempts to commission its cross-connect to ensure the ODR SPVC traverses this new link as it is signalled to the next node in the primary path. If this re-route fails, the connection is cranked back to the source node. At this point the source node waits an interval of time before re-routing the ODR SPVC.

[0039] After the waiting interval elapses, the source node re-initiates the ODR SPVC using the sequence of events described above. Successive attempts to re-

route the ODR SPVC using the primary path will result in the source node increasing the interval of time before repeating this sequence. The source node strives to setup the ODR SPVC along the nodes and specified links in its primary path until it is successful.

[0040] It should be noted that a node failure along the specified path results in the source node attempting to re-route the ODR SPVC along its primary path. However, the re-route will not succeed unless the failed node is restored. The network nodes are unable to re-route the ODR SPVC on nodes not explicitly specified by the operator because this re-routing scheme restricts the ODR SPVC to routes along the nodes in the primary path. A link failure along this path results in the source node re-routing the ODR SPVC along its primary path. However, the success of the re-route depends on the configuration of the primary path of the ODR SPVC; the configuration of the network in the area of the failure; and the status of the failed link at the time of the re-route.

[0041] The re-route may succeed if the operator explicitly selected the failed link as part of the primary path and the failed link is restored when the appropriate node re-routes the ODR SPVC over the link. If the link is not restored and parallel links exist between the adjacent nodes, the node will not re-route the ODR SPVC around this link failure because this re-routing scheme restricts the ODR SPVC to the links specified by the operator.

[0042] The re-route may succeed if the link is the sole link between adjacent nodes in the primary path and it is restored when the appropriate node re-routes the ODR SPVC over the link. In this case it does not matter whether the operator explicitly selected the link as part of the primary path. Since this link is the only link between the adjacent nodes (i.e., parallel links do not exist between the adjacent nodes), the re-route will not succeed until the failed link is restored.

[0043] The re-route may succeed if the operator did not select the failed link as part of the primary path and parallel links exist between the adjacent nodes. If the failed link is not restored, the appropriate node's call control software will select a link other than the failed link and attempt to commission its cross connects to ensure the ODR SPVC traverses this new link as it is signalled to the next node in the primary path.

• Primary Path-Alternate Path

[0044] Under the primary path-alternate path re-routing scheme, an ODR SPVC is restricted to being re-routed only along the nodes and specified links in its primary path or alternate path. Since the primary path is the preferred path of an ODR SPVC, the source node will attempt to set up the ODR SPVC on this path first. Failure to route the ODR SPVC along the nodes and links in this path will result in the crankback of the connection to its source node. At this point the source node

either immediately re-routes the connection along the nodes and links in its primary path or immediately re-routes the connection along the nodes and links in its alternate path.

[0045] If the failure is due to a link not specified by the operator, the source node immediately re-routes the ODR SPVC along its primary path. This is done to give the network nodes another opportunity to immediately set up the connection along the nodes and links in the primary path whenever parallel links exist in the area of the link failure. If parallel links exist in the area of the link failure, the appropriate node's call processing means selects a link other than the failed link. The node attempts to commission its cross connect to ensure the ODR SPVC traverses this new link as it is signalled to the next node in the primary path. If the re-route fails, the connection is cranked back to the source node. At this point the source node immediately re-routes the connection along the nodes and links in its alternate path.

[0046] Failure to re-route the ODR SPVC along the nodes and links in its alternate path will result in the crankback of the connection to its source node. At this point the source node either immediately re-routes the ODR SPVC along the nodes and links in its alternate path or waits an interval of time before re-routing the ODR SPVC. If the failure is due to a link not specified by the operator, the source node immediately re-routes the ODR SPVC along its alternate path. This is done to give the network elements another opportunity to immediately set up the connection along the nodes and links in the alternate path wherever parallel links exist in the area of the link failure. If parallel links exist in the area of the link failure, the appropriate node's call control software selects a link other than the failed link. The node attempts to commission its cross connect to ensure the ODR SPVC traverses this new link as it is signalled to the next node in the alternate path. If this re-route fails, the connection is cranked back to the source node. At this point the source node waits an interval of time before re-routing the ODR SPVC.

[0047] After the waiting interval elapses, the source node re-routes the ODR SPVC using the sequence of events described above. Successive attempts to re-route the ODR SPVC using the primary path-alternate path sequence will result in the source node increasing the interval of time before repeating the sequence. The source node strives to setup the connection along the nodes and links in its primary path or alternate path until it is successful.

• Primary Path-Any Path

[0048] Under the primary path-any path re-routing scheme, the source node attempts to route an ODR SPVC along the nodes and links in its primary path first. Failure to route the ODR SPVC along this path will result in the crankback of the connection to its source

node. At this point the source node either immediately re-routes the ODR SPVC along the nodes and links in its primary path or immediately re-routes the ODR SPVC using conventional SPVC routing (i.e., using SVC routing).

[0049] If the failure is due to a link not specified by the operator, the source node immediately re-routes the ODR SPVC along its primary path. This is done to give the network elements another opportunity to immediately set up the connection along the nodes and links in the primary path whenever parallel links exist in the area of the link failure. If parallel links exist in the area of the link failure, the appropriate node's call control software selects a link other than the failed link. The node attempts to commission its cross connect to ensure the ODR SPVC traverses this new link as it is signalled to the next node in the primary path. If the re-route fails, the connection is cranked back to the source node. At this point the source node immediately re-routes the ODR SPVC via conventional SPVC or SVC routing.

[0050] If SVC routing fails, the source node waits an interval of time before re-routing the ODR SPVC. After the time elapses, the source node re-routes the ODR SPVC using the sequence of events described above. Successive attempts to re-route the ODR SPVC using the primary path-any path sequence will result in the source node increasing the interval of time before repeating the sequence. The source node strives to setup the connection until it is successful.

• Primary Path-Alternate Path-Any Path

[0051] Under the primary path-alternate path-any path re-routing scheme, the source node will attempt to set up the ODR SPVC on the primary path first. Failure to route the ODR SPVC along the nodes and links in this path will result in the crankback of the connection to its source node. At this point the source node either immediately re-routes the connection along the nodes and links in its primary path or immediately re-routes the connection along the nodes and links in its alternate path.

[0052] If the failure is due to a link not specified by the operator, the source node immediately re-routes the ODR SPVC along its primary path. This is done to give the network nodes another opportunity to immediately set up the connection along the nodes and links in the primary path whenever parallel links exist in the area of the link failure. If parallel links exist in the area of the link failure, the call processing means of appropriate node selects a link other than the failed link. The node attempts to commission its cross connect to ensure the ODR SPVC traverses this new link as it is signalled to the next node in the primary path. If the re-route fails, the connection is cranked back to the source node. At this point the source node immediately re-routes the connection along the nodes and links in its alternate path.

[0053] Failure to re-route the ODR SPVC along the nodes and links in its alternate path will result in the crankback of the connection to its source node. At this point the source node either immediately re-routes the ODR SPVC along the nodes and links in its alternate path or immediately re-routes the ODR SPVC via SVC routing. If the failure is due to a link not specified by the operator, the source node immediately re-routes the ODR SPVC along its alternate path. This is done to give the network nodes another opportunity to immediately set up the connection along the nodes and links in the alternate path whenever parallel links exist in the area of the link failure. If parallel links exist in the area of the link failure, the appropriate node's call processing means selects a link other than the failed link. The node attempts to commission its cross connect to ensure the ODR SPVC traverses this new link as it is signalled to the next node in the alternate path. If the re-route fails, the connection is crankback to the source code. At this point the source node immediately re-routes the ODR SPVC via conventional SPVC or SVC routing.

[0054] If SVC routing fails, the source node waits an interval of time before re-routing the ODR SPVC. After the waiting interval elapses, the source node re-routes the ODR SPVC using the sequence of events described above. Successive attempts to re-route the ODR SPVC using the primary path-alternate path-any path sequence will result in the source node increasing the interval of time before repeating the sequence. The source node strives to set up the connection until it is successful.

[0055] Fig. 5 illustrates a preferred relational database structure 52 used by a given network node 32 to keep track of ODR SPVCs managed thereby. The node database structure includes the following tables: (a) an indexed table 54 of P-NNI network node identifiers; (b) an indexed table 56 of "compressed" ODR SPVC DTLs, as explained in greater detail below, and an ODR SPVC list 58, part 59 of which is stored in random access memory. The ODR SPVC database includes one record for each ODR SPVC which has originated from or is managed by the node. Each ODR SPVC record includes:

- (a) an "operatorDrtRtng" field 60 which specifies whether the corresponding ODR SPVC is a conventional SPVC or an ODR SPVC;
- (b) a "prnmNtwrkPathIndex" field 62 which points to an entry in the compressed DTL table that represents the primary path of an ODR SPVC;
- (c) a "altNtwrkPathIndex" field 64 which points to the alternate path entry in the compressed DTL table for the ODR SPVC; and
- (d) a "re-route scheme" field 66 which stores the re-route scheme for the ODR SPVC.

[0056] In this manner, the source node of an SPVC can determine whether the SPVC is an ODR SPVC, and, if so, determine the attributes associated with the ODR SPVC in order to take appropriate action in the event of link failure.

[0057] The compressed DTL table 54 has a "compressed DTL" field 68 for storing network paths in a compressed format. A compressed network path is illustrated in greater detail at reference no. 68' and comprises a sequence of link identifiers (i.e. P-NNIPortId-field 70), and pointers 72 to the node table 54. In the preferred embodiment, P-NNIPortId=0 signifies that no link has been specified by the operator, whereby intermediate nodes in the path of the ODR SPVC are free to select whichever link is available on the node.

[0058] Those skilled in the art will appreciate that in the P-NNI routing protocol a "node" can in fact represent an entire subnetwork or "peer group". For example, a node identifier may be a sub-prefix of an ATM node address (itself a prefix of an end system ATM address) which summarizes the reachability of all nodes in the subnetwork or peer group. The NMS 46, however, is typically connected to each node in an entire network, including the nodes in the subnetworks thereof, so that the operator can be made aware of and manually select each physical switching element, and optionally, the links, in each subnetwork along the path of an ODR SPVC. In this case, the DTL in SPVC Call Setup message will specify the complete path (i.e., at least each and every physical switch) to the SPVC destination end point. In this embodiment, since the complete path is specified in the DTL, in the event of a link failure, the crankback procedures discussed previously operate to crank back the connection request back to the originating source node.

[0059] Alternatively, however, the operator may specify a subnetwork in the path of the ODR SPVC, without detailing the specific nodes to be traversed in the subnetwork. In this case, the node at the entry to the peer group or subnetwork (hereinafter the "lead node") computes the path through the subnetwork to a following node or subnetwork specified in the DTL of the SPVC Call Setup message. In such an alternative embodiment, in the event of a link failure in the subnetwork, the aforementioned P-NNI crankback procedures operate to crankback the connection request to the lead node of subnetwork (since it is a DTL originator), whereby the lead node may compute an alternate path through the subnetwork to reach the following node or subnetwork specified in the DTL of the connection request. If the reroute is unsuccessful, the lead node then cranks the connection request back to the originating source node so that it may apply the ODR SPVC reroute scheme.

[0060] The preferred embodiment has described an ODR connection for an SPVC. Those skilled in the art will appreciate that the invention may alternatively be applied in relation to SVCs, wherein the ingress network node transmits the ODR source-routed call setup mes-

sage across the network. Alternatively, the user interface means may communicate directly with the end-station and the SVC call setup message may carry the manually routed connection path(s) and re-routing restrictions. It will also be understood that the invention may be applied in relation to SVC or SPVC point-to-multipoint and multipoint-to-point connections, wherein each leaf is separately set up using a separate connection request message. Furthermore those skilled in the art will appreciate that the invention may be applied in relation to router based (i.e. layer 3 switching) network environments such as tag switching or multi-protocol label switching (MPLS) networks which feature source routed connection setups. Similarly, those skilled in the art will appreciate that numerous other modifications and variations may be made to the preferred embodiments described above without departing from the spirit and scope of the invention.

Claims

1. A method of establishing a connection in a communications network comprising interconnected network elements, said method comprising:
 - (a) providing a predetermined preferred path definition to a source network element, said path definition comprising a destination network element and intermediate network elements between said source and destination network elements;
 - (b) creating a source-routed connection request message which specifies said preferred path; and
 - (c) signalling said connection request message from said source network element to said destination network element along said preferred path in order to establish said connection.
2. The method according to claim 1, wherein said network element comprises at least one of:
 - a subnetwork;
 - a network node; and
 - an end-station.
3. The method according to claim 2, wherein said predetermined path is manually provisioned.
4. The method according to claim 3, including establishing bearer channel cross-connection on at least each said intermediate network element in order to connect adjacent network elements along said preferred path.
5. The method according to claim 4, including the optional steps of provisioning a link in said preferred path definition and specifying the link in the

source-routed connection request message.

6. The method according to claim 5, wherein at least said intermediate network elements are nodes, each of which commissions a bearer channel cross-connection thereon so as to use those links specified in the source-routed connection request message.
7. The method according to claim 6, wherein adjacent network elements automatically select a link in the event the source-routed connection request message does not identify the link to be used between adjacent elements.
8. The method according to claim 1, including manually selecting a re-routing scheme and storing the re-routing scheme on the source network element.
9. The method according to claim 8, including manually provisioning one or more alternate paths of network elements and optional links for the connection, and storing the alternate paths on the source network element.
10. The method according to claim 9, wherein the re-routing scheme constrains the path along which the connection may be re-routed in the event of an inoperative link along the preferred path, said constraint being one of (i) the preferred path only; (ii) the alternative paths; (iii) any possible path from the source network element to the destination network element; and (iv) the alternative paths, and if none of the alternative paths are available, any possible path from the source to the destination network elements.
11. The method according to claim 10, wherein the network is a P-NNI network and the source-routed connection message is one of an SPVC call setup message and an SVC call setup message which include a designated transit list (DTL) specifying each intermediate network element between the source and destination network elements.
12. The method according to claim 11, including cranking back the connection request to the source network element for re-routing the connection in the event the link becomes inoperative during the establishment of the connection.
13. The method according to claim 11, including signalling the source network element to re-route the connection in the event the link becomes inoperative after the establishment of the connection.
14. The method according to claim 3, wherein said manual provisioning includes one of:

manually entry;
manual with point and click assistance; and
automatic route generation with optional manual editing.

15. A method of establishing a connection in a sourced-routed, connection-orientated, communications network which comprises a signalling protocol for forwarding connection requests to network nodes, said method comprising:

(a) manually provisioning a preferred path for the connection, including a source node, a destination node, and intermediate nodes or sub-networks therebetween;
(b) manually provisioning a re-routing restriction for the connection;
(c) creating a source-routed connection request message which specifies the preferred path;
(d) each node or subnetwork along the preferred path successively signalling the connection request message to a following node or subnetwork along the preferred path and commissioning a bearer channel cross-connection on each said specified node, or on nodes located in said subnetwork which have been identified by the subnetworks themselves, in order to connect adjacent nodes or subnetworks along the preferred path; and
(e) signalling the source node in the event of a blockage in the preferred path and, in such event, re-initiating steps (c) and (d) using a path permitted by the manually provisioned re-routing restriction.

16. A communications network, comprising:

user interface means for manually provisioning a network path for a connection and specifying at least a source network element, a destination network element, and each intermediate network element therebetween;
a plurality of interconnected network elements, each comprising call processing means for receiving a source-routed connection request message incorporating a transit list, establishing a bearer channel cross-connection linking the network element to a previous network element and a following network element in the transit list, and forwarding the connection request message to the following network element in the transit list,
wherein the user interface means is connected to the source network element and commands the source network element to create and signal a source routed connection request message having a transit list specifying the

manually provisioned network path.

17. A communications network, comprising:

user interface means for manually provisioning a preferred network path for a connection and a rerouting scheme therefor, said preferred network path specifying at least a source network element, a destination network element, and each intermediate network element therebetween;
a plurality of interconnected network network elements, each comprising call processing means for (a) receiving a source-routed connection request message incorporating a transit list, (b) establishing a bearer channel cross-connection linking the network element to a previous network element and a following network element in the transit list, (c) forwarding the connection request message to the following network element in the transit list, and (d) cranking back the connection request to the source network element in the event the preferred network path is blocked and re-routing the connection in accordance with the re-routing scheme,
wherein the user interface means is connected to the source network element and signals the source network element to create and signal a source routed connection request message having its transit list specifying the manually provisioned network path.

18. The network according to claim 17, wherein the rerouting scheme comprises one of (a) the preferred network path only, (b) at least one alternate manually provisioned network path, and (c) any available network path to the destination network element.




Figure 1




Figure 2A




Figure 2B

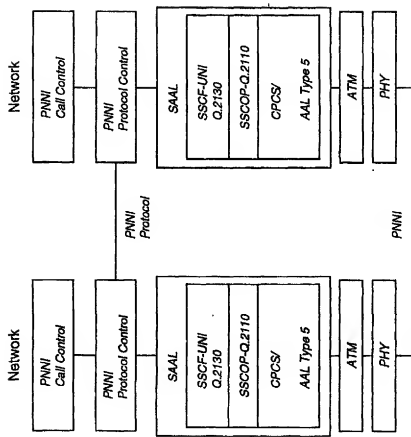



Figure 4






Figure 5